

Testing and Improving Features in PINT

JOSEF ZIMMERMAN^a ¹ AND SCOTT RANSOM ²

¹*University of Virginia
Department of Physics
382 McCormick Road*

Charlottesville, VA 22904, USA

²*National Radio Astronomy Observatory
520 Edgemont Road
Charlottesville, VA 22903, USA*

(Received August, 2020)

ABSTRACT

PINT is a new pulsar timing software designed in `python`. A model for the effects of the troposphere delay was added to PINT, accounting for the hydrostatic component of the delay. The implementation utilizes the rigorously tested `Astropy` library, ensuring that the calculations are reliable and vectorized for efficiency. Additional testing was performed on the PINT model for planetary Shapiro delay, resulting in the discovery of a bug in `Tempo2`'s model that had gone uncaught for nearly 15 years. Both models were assessed for importance in the NANOGrav search for gravitational waves.

Keywords: pulsars, pulsar timing

1. INTRODUCTION

PINT¹ (PINT is not `Tempo3`, Luo et al. (Forthcoming)) is a novel pulsar timing library designed to improve upon the previous `TEMPO`² and `Tempo2`³ (Hobbs et al. 2006). As a pulsar timing package, PINT allows a user to process a large number of radio pulse Time of Arrival (TOA) measurements, making corrections based on timing conventions, then utilizes various fitting algorithms to construct a timing model for an individual pulsar. That model predicts the exact time of expected pulse arrival to within 10 ns precision. Ultimately, by constructing a pulsar timing array and accounting for all predicted effects that cause variation in pulse arrival time, the residual difference between the modeled and observed TOAs are recovered, allowing one to search for unmodeled gravitational wave signals. With such gravitational wave search efforts well underway (Arzoumanian et al. 2018), it is of the utmost importance to ensure that the pulsar timing array sensitivity is limited by the number of pulsar observations and inherent noise, rather than errors in computational procedures. Thus, the primary

goal of PINT is to provide an independent implementation of pulsar timing algorithms, using rigorously tested scientific packages such as `Astropy` (Price-Whelan et al. 2018) whenever possible, in order to cross check the predictions of `TEMPO` and `Tempo2`.

As a secondary goal, PINT is developed in `python` using a highly modular design, allowing a user to easily add new functionality themselves to conduct any necessary science. Whereas `TEMPO` and `Tempo2` were developed in `FORTRAN` and `C`, respectively, making their procedures very difficult to modify and recompile. Contrasting this, `python` is one of the most popular programming languages today, making it easily accessible to most astronomers.

One feature that PINT previously lacked was a method to model pulse delay due to effects of the troposphere. For any topocentric TOAs, pulses are delayed on the order of 10 ns at zenith (Hobbs et al. 2006), increasing significantly for observations closer to the horizon. As `Tempo2` already accounts for this effect, this paper details the steps taken to modify PINT to appropriately model the troposphere delay. The troposphere models of PINT and `Tempo2` were compared to check for consistency, both in terms of the algorithmic implementation and correct agreement with the expected model behavior. Additional testing of PINT was performed, com-

¹ <https://github.com/nanograv/PINT>

² <https://sourceforge.net/projects/tempo/>

³ <https://bitbucket.org/psrsoft/tempo2/src/master/>

paring predictions of the planetary Shapiro delay to the models of `Tempo2`.

Finally, tests were performed on the NANOGrav 12.5 year data set to quantify the importance of the troposphere and planetary Shapiro delays in aiding the search for gravitational waves. The magnitude and effect of not modelling each delay was assessed and compared to the target sensitivity levels necessary to detect gravitational waves. This effort was intended to determine if either effect is negligible enough to justify its omission in order to decrease the computational requirements for processing data in a pulsar timing array.

2. TROPOSPHERE DELAY

The troposphere delay is a small effect causing radio pulses to arrive slightly later after being scattered by molecules in the atmosphere. As conventionally modeled in very long baseline interferometry (VLBI), the troposphere delay is considered the sum of the “hydrostatic” and “wet” components. The hydrostatic delay is the native delay from the structure of the atmosphere, and at any given site remains nearly constant aside from annual variations. Importantly, it can be empirically modeled for any site without any knowledge of the local weather conditions on any given day of observation. Conversely, the wet delay encompasses all of the daily variation, primarily emerging from changes to the water vapor content of the atmosphere. As it is much more dynamic, and requires additional data from the observatory, it is much more difficult to model. As implemented in `Tempo2`, the user must provide their own estimates of the wet delay at zenith for each TOA in order to be modeled properly. Due to this hurdle of accessibility, in addition to the relative smallness of the wet delay compared to the hydrostatic delay (10x smaller at zenith, [Hobbs et al. \(2006\)](#)), the PINT troposphere model currently only accounts for the hydrostatic component.

The first step in modeling the troposphere delay is to calculate the effect at zenith Δ_z . This is accomplished using the often-referenced analytical model of [Davis et al. \(1985\)](#) (used by `Tempo2`), which accounts for corrections based on the centrifugal acceleration on the surface and the nonlinearity of the atmosphere at higher elevations:

$$\Delta_z = \frac{P/(43.921\text{kPa})}{c(1 - 0.00266 \cos(2\phi) - 2.8 \times 10^{-4}H/\text{km})}, \quad (1)$$

where P is the pressure at the observatory, c is the speed of light, ϕ is the latitude of the observatory, and H is the height of the observatory above the geoid.

As implemented in `Tempo2`, the user must specify the pressure at the observatory, otherwise it will assume

STP conditions of 101.325 kPa. As most radio telescopes are located at significantly higher elevations than sea level, this default approximation is not very useful. To increase the ease of use for any user, the PINT troposphere model automatically approximates the pressure at the observatory by using an approximate model for pressure variation with elevation: ([CRC Handbook 2004](#))

$$Z \equiv \frac{R_{\oplus} H}{R_{\oplus} + H} \quad (2)$$

$$T = (288.15 - 0.0065 * H/\text{m})\text{K} \quad (3)$$

$$P = 101.325\text{kPa} \times (288.15\text{K}/T)^{-5.25575}, \quad (4)$$

where Z is the geopotential height and T is the linear correction for atmospheric lapse rate. For small elevations (< 11 km, [CRC Handbook \(2004\)](#)), this formula accurately describes the pressure behavior for increasing altitude. PINT uses the internally stored locations and elevations of all radio observatories to automatically calculate these quantities with ease for any user.

The final step in calculating the troposphere delay is to utilize a mapping function to determine the delay at any altitude ϵ in the sky, rather than just at zenith. The first order approximation to this assumes a plane-parallel atmosphere, yielding a mapping function $m(\epsilon) = 1/\sin \epsilon$, representing how radio pulses arriving near the horizon travel through significantly more atmosphere than those originating at the zenith. A more accurate representation is obtained by the Niell Mapping Function (NMF, [Niell \(1996\)](#)), which modifies the plane parallel model to include corrections for both the curvature of the atmosphere and seasonal variations:

$$m(\epsilon) = \frac{1 + \frac{a}{1 + \frac{b}{1+c}}}{\sin \epsilon + \frac{a}{\sin \epsilon + \frac{b}{\sin \epsilon + c}}}. \quad (5)$$

The definitions of the corrective terms a , b , and c are defined in Equation 5 and Table 3 of [Niell \(1996\)](#). Each term varies with the latitude of the observatory and fraction of year at the time of observation, which are both calculated by the PINT implementation automatically. The most important step in the mapping function is determining the altitude ϵ of the pulsar at the time of observation, which is unique for each TOA and depends on both the latitude and longitude of the observatory. This calculation is preformed in an efficient and reliable method using the `Astropy` vectorized libraries, ensuring that it is bug-free and computationally fast enough to include in most modeling situations.

To test the PINT implementation of the troposphere delay, the low declination pulsar J1909-3744 observations from the GBT in the 12.5 year data set were considered. For each TOA, the troposphere delays from

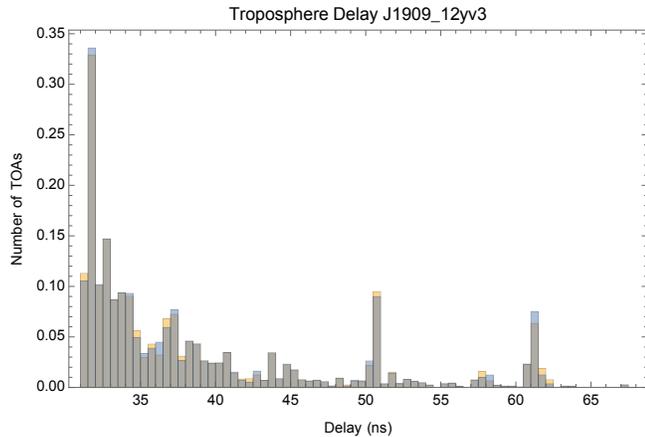


Figure 1. Comparison of troposphere delay calculations by PINT and Tempo2 for TOAs observed by the GBT for J1909 show very good agreement. PINT delays are shown in blue, while Tempo2 is shown in orange. Most of the plot appears grey from the overlap of the two histograms. Note that for this comparison with Tempo2, the pressure scaling function of PINT was not used.

PINT and from Tempo2 using the `general2` plugin were calculated, generating a histogram showing the distribution of delays in Fig. 1. The results show very good agreement between the two software models.

After verifying the accuracy of the PINT troposphere model, the NANOGrav 12.5 year data set was analyzed to determine the effect of including troposphere modeling on overall accuracy of the timing model. For each pulsar observed from the GBT, the general least squares (GLS) fitting routine of PINT was used, comparing a fit with the troposphere effects modeled to a fit without the modeling enabled. For each TOA, the difference between the observed arrival time and predicted model was calculated as the residual. Then, the difference between the residual with the troposphere model and the same residual without the troposphere model was calculated for each TOA. The resulting data set was plotted as a histogram (Fig. 2), with positive values indicating that modeling accuracy for that TOA improved with the inclusion of the troposphere model, and negative indicating worsening. On average, inclusion of the troposphere models improves the TOA residuals by 0.692 ± 0.060 ns, a small but significant improvement. Because the residual improvement is so small relative to the order of magnitude of the troposphere effect, the χ^2 fitter is most likely overcompensating for the effect by tweaking other models, making those inaccurate, when the troposphere delay is not modeled properly. This indicates that in addition to slightly increasing the residual accuracy of the PINT model, inclusion of the troposphere delay increases the accuracy of other model fit parameters.

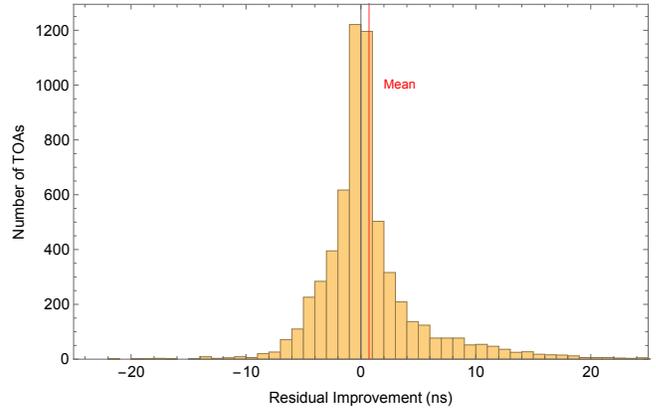


Figure 2. This histogram shows the distribution of change in residuals with and without the troposphere model enabled for TOAs observed by the GBT. Positive numbers represent improvement of the residuals by enabling troposphere modeling. Assuming the set is randomly sampled, the mean and $1\text{-}\sigma$ uncertainty for average improvement is 0.692 ± 0.060 ns, shown as a red shaded region on the graph.

In the future, new features may be added to the troposphere delay model in PINT, including the ability to handle wet delays and specify actual atmospheric pressures for each TOA.

3. SHAPIRO DELAY

The second model of PINT investigated in this paper is the planetary Shapiro delay (Shapiro 1964). As this paper only considers tests involving the Shapiro delay, rather than implementing the model itself, the mathematical details will not be discussed.

The Shapiro delay occurs when a radio pulse travels through the gravitational well of a massive object. The strength of the delay increases significantly for more massive objects, and objects appearing very near to the target pulsar in the sky. By default, PINT includes Shapiro delay calculations for the Sun in all models; however, while it has the ability to do the same calculation for each of the planets, it is an optional parameter disabled by default.

To test PINT's ability to model the planetary Shapiro delay, simulated TOAs were generated over a 20 year period for a pulsar in NGC 6440, a globular cluster near the ecliptic. As each of the planets orbit the Sun, they appear to move along the ecliptic, thus bringing them in close proximity to NGC 6440 once per revolution, which should cause a spike in the Shapiro delay. The predicted Shapiro delay was calculated by PINT and plotted in Fig. 3 as a function of angular separation between each planet and the target pulsar, correctly showing that the largest delay occurs when the separation is small.

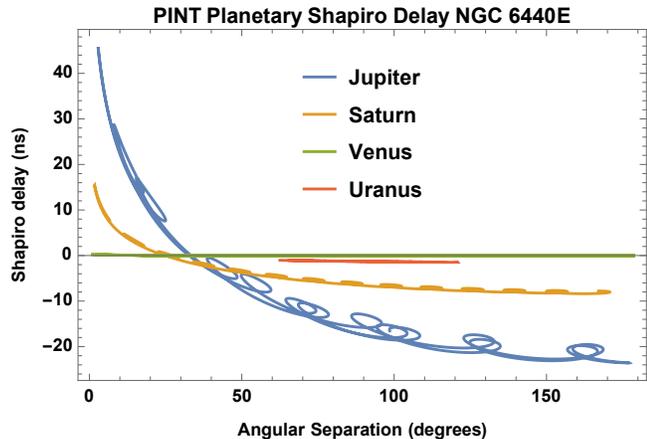


Figure 3. Planetary Shapiro delays calculated by PINT over a 20 year period show expected behavior. The maximum delay from each planet occurs when it is nearest in alignment to the target pulsar in NGC 6400. The smaller looping structure is a result of annual variations from Earth’s motion relative to each of the planets.

The same calculation was performed in Tempo2, yielding a different result shown in Fig. 4. Instead of the expected behavior, according to Tempo2, the Shapiro delay is maximized when each planet is on the opposite side of the sky from NGC 6440. This is a bug in the code of Tempo2; while the discrepancy had been noticed early on in the development of PINT, it was never followed up on. The bug has been reported to the development team of Tempo2, who acknowledged it and is working on correcting it. Additionally, the Parkes Pulsar Timing Array (PPTA) is reassessing their previous analysis in the search for gravitational waves, which relies on the planetary Shapiro delay calculation from Tempo2.

After verifying that the planetary Shapiro delay calculation is handled properly by PINT, tests were performed to test its importance in the search for gravitational waves. For each NANOGrav pulsar, 20 years of observation were simulated, calculating the sum of delays for each planet. Figure 5 shows the total Shapiro delay for the pulsar J1614–2230, which demonstrated the greatest variation of all NANOGrav pulsars over the simulated observation time period, an 80 ns effect. Highlighted in that image is an orange sinusoid, fit to the Shapiro delays using a least squares regression, naively representing how the Shapiro delay can be confused with a gravitational wave signal due to its large amplitude and long period. Although actual gravitational wave search algorithms are more complicated than performing a regression fit on a single pulsar, this demonstrates the necessity of properly modeling the planetary Shapiro delay to maximize the precision of a pulsar timing array.

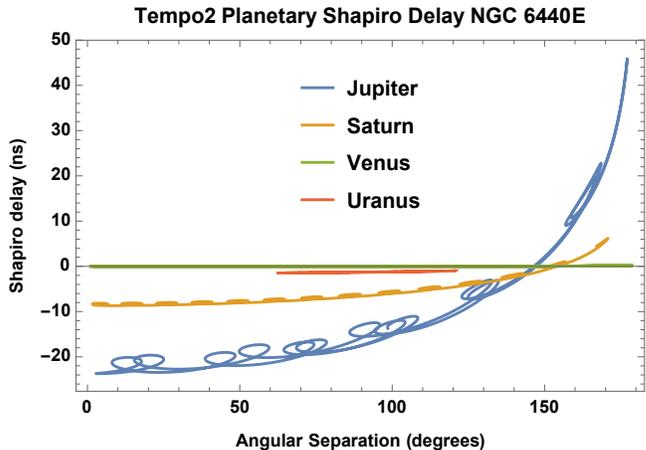


Figure 4. Planetary Shapiro delays calculated by Tempo2 over a 20 year period show unexpected behavior, with the maximum Shapiro delay occurring as each planet nears 180° separation from NGC 6440. This is the result of a bug in Tempo2 miscalculating the Earth-Jupiter vector direction.

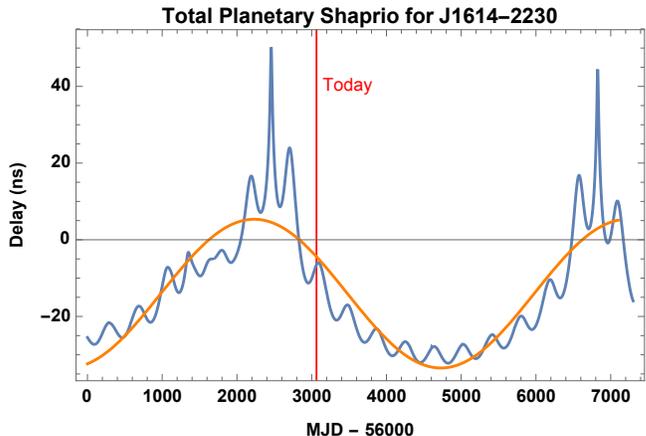


Figure 5. The total planetary Shapiro delay as modeled by PINT for the pulsar J1614 shows the largest variation of 80 ns over a 20 year period of simulated observations. Superimposed in orange is a sinusoidal fit to the delay, demonstrating possible confusion of the effect with a 12 year period, 40 ns amplitude gravitational wave signal. Failing to properly account for the planetary Shapiro delay may significantly impede the search for gravitational waves.

Lastly, Fig. 6 shows a histogram of the distribution of Shapiro delay variations for every pulsar used by NANOGrav. Note that every single pulsar has a variation over 10 ns, even when far away from the ecliptic. Thus, in order to utilize PINT to its greatest potential, the planetary Shapiro delay modeling should be enabled for gravitational wave searches.

4. ONGOING WORK

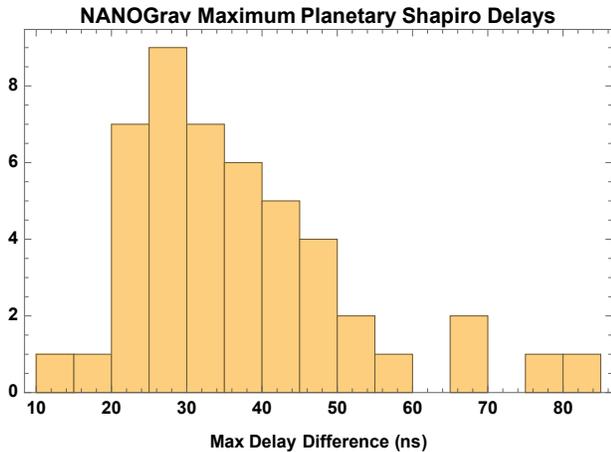


Figure 6. The distribution of planetary Shapiro delay variations for each NANOGrav pulsar over 20 years of simulated observation shows a significant (> 10 ns) impact of the delay.

Current work is ongoing to perform bug fixes in PINT and continue the investigation of the planetary Shapiro delay. There appears to be an inconsistency in PINT’s modeling of solar wind dispersion effects, causing the effects to appear out of phase by half of a year. While the cause of the problem has not been identified yet, investigation is continuing.

A more rigorous analysis of the planetary Shapiro delay is being performed with the goal of quantifying its similarity to a false gravitational wave signal. Gravitational wave detectability is most often quantified by the Hellings and Downs curve, which predicts the correlation of pulse arrival time correlations between two pulsars around the sky (Hellings & Downs 1983):

$$C = \frac{1 - \cos \gamma}{2} \ln \left(\frac{1 - \cos \gamma}{2} \right) - \frac{1}{6} \frac{1 - \cos \gamma}{2} + \frac{1}{3}. \quad (6)$$

The curve is calculated in terms of the angular separation γ , and is maximized for small values of γ , reaches a minimum near 90° , and increases again for pulsars on

opposite sides of the sky. Current investigation focuses on numerically creating a similar curve for the effect of the planetary Shapiro delay, which should be most correlated at small γ and have negative correlation for large values of γ . This curve will be recreated by pairing all combinations of two NANOGrav pulsars over a large period of observations, and calculating the correlation as a function of the angular separation. Comparison of such a plot to the Hellings and Downs curve will further quantify the importance of modeling the planetary Shapiro delay.

5. CONCLUSION

This paper summarizes the work done to implement and test a troposphere delay model in PINT. A model for the hydrostatic component of the troposphere delay was implemented into PINT, and compared to the corresponding model in *Tempo2* to validate both models. Inclusion of the troposphere model in fitting routines makes minor but significant increases the model accuracy by a small amount, likely removing systematic errors in the measurement of all other parameters in the fitted model. Additional tests were performed on the planetary Shapiro delay model of PINT and *Tempo2*, discovering a bug in *Tempo2* that had gone uncorrected for nearly 15 years and prompting a revised analysis of the results of the PPTA. Simulated observations of PSR J1614-2230 showed a planetary Shapiro delay variation of approximately 80 ns over 20 years. Importantly, the delay time dependence appears similar in amplitude and period to the effect of a gravitational wave signal, indicating that it may be confused with a gravitational wave signal if not modeled correctly. Ultimately, regardless of the proximity to the ecliptic, all NANOGrav pulsars demonstrate significant variation due to the Shapiro delay, and modeling should be enabled and properly handled by PINT in future gravitational wave searches to maximize sensitivity.

REFERENCES

- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *The Astrophysical Journal Supplement Series*, 235, 37, doi: [10.3847/1538-4365/aab5b0](https://doi.org/10.3847/1538-4365/aab5b0)
- CRC Handbook. 2004, *CRC Handbook of Chemistry and Physics*, 85th Edition, 85th edn., ed. D. R. Lide (CRC Press)
- Davis, J. L., Herring, T. A., Shapiro, I. I., Rogers, A. E. E., & Elgered, G. 1985, *Radio Science*, 20, 1593, doi: [10.1029/RS020i006p01593](https://doi.org/10.1029/RS020i006p01593)
- Hellings, R. W., & Downs, G. S. 1983, *ApJL*, 265, L39, doi: [10.1086/183954](https://doi.org/10.1086/183954)
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, 369, 655, doi: [10.1111/j.1365-2966.2006.10302.x](https://doi.org/10.1111/j.1365-2966.2006.10302.x)
- Luo, J., Ransom, S., Demorest, P., et al. Forthcoming
- Niell, A. E. 1996, *Journal of Geophysical Research: Solid Earth*, 101, 3227, doi: [10.1029/95JB03048](https://doi.org/10.1029/95JB03048)
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *The Astronomical Journal*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)

Shapiro, I. I. 1964, Phys. Rev. Lett., 13, 789,
doi: [10.1103/PhysRevLett.13.789](https://doi.org/10.1103/PhysRevLett.13.789)